arXiv:1610.07080v1 [cs.FL] 22 Oct 2016

An Alternating Automaton for First-Order Linear Temporal Logic

Tech Report

Yannick Lebrun, Raphaël Khoury, and Sylvain Hallé

Laboratoire d'informatique formelle
Université du Québec à Chicoutimi, Canada

Abstract. In this paper we give automata-based representation of LTL-FO$^+$ properties. LTL-FO$^+$ is an extension of LTL that includes first-order quantification over bounded variable, thus greatly increasing the expressivity of the language. An automata representation of this formalism allows greater ease in writing and understanding properties, as well as in performing manipulations, such as negation or emptiness checking. The automata representation of an LTL-FO$^+$ formula has finite size regardless of the domain of quantified variables, and the number of states that is linear in the size of the property.

1 Introduction

LTL-FO$^+$ [6] is a formal language used for the specification of trace properties which distinguishes itself from other representations by its exceptional expressiveness. It allows users to state finer relationships between the different elements of several messages in complex event trace. For example, in [6] LTL-FO$^+$ is used to express properties related to XML messages traces generated by web services. Such properties cannot be stated with a less expressive formalism such as LTL.

However, the on-the-fly verification algorithm is based on the decomposition and rewriting of formulæ. It is highly intensive in its space consumption, with multiple manipulation being performed even when processing messages that have no bearing on the validity of the formula. Indeed, for some formulæ, the evaluation tree expands indefinitely, and equivalent subtrees, which could be pruned, are hard to identify. Furthermore, the elaborate syntax of LTL-FO$^+$ can make it difficult to state and read properties (see for example [7]).

An automat-based representation of LTL-FO$^+$ formulæ would thus ease both the writing and reading of formulæ. The space and time overhead of the verification process would also be optimized, since the validity of a formula could be ascertained by maintaining a list of current valuations of and current state or states. Furthermore, an automata representation will allow multiple useful manipulations to be performed with ease, notably counterexamples generation, emptiness checking, intersection and negation of properties.

The automata we propose is a variation of Vardi's alternating automata [10], which we have enriched with first-order quantifiers over a finite set of formula

variables. This makes it easier to express intricate formulæ over complex events, where each event consists of a XML object with possibly multiple valuations for each path. Like the alternating automata, the proposed automata distinguishes between *existential* and *universal* transitions. Existential transitions are analogous to non-deterministic transitions in regular Büchi automata. Upon encountering such a transition, the automata can be though of as choosing between multiple destination states. Conversely, when encountering a universal transition, the automata continues its run in both target states simultaneously. Because of the presence of universal and existential quantifiers, a run over an alternating automata generates a tree of states. A run is accepting of there *exists* at least one tree for which *every* branch visits an accepting state infinitely often. While alternating automata are equally expressive as non-deterministic Büchi automata, we show in this paper how the notions of existential and universal transition can be used to model the quantifiers present in LTL-FO$^+$ formulæ. The automata is additionally enriched with a partial function mapping formula variables to variables to their values. This function is manipulated by the automata's transition as the input sequence is read and consulted to determine the truth value of elementary propositions.

In this paper, we show how to extend Vardi's alternating automata to accommodate the greater expressivity of LTL-FO$^+$. Section 2 provides background information about LTL-FO$^+$. Section 3 surveys existing automata representations for other formal logics. In Section 4, we show how to construct an modified alternating automata from an LTL-FO$^+$ formula, such that the automata accepts exactly the same set of input sequences as the original property. Section 5 sketches out a proof of correctness. Concluding remarks are given in Section 6.

2 The First-Order Temporal Logic LTL-FO$^+$

LTL-FO$^+$, a first-order extension of a well-known logic called Linear Temporal Logic (LTL); LTL has already been suggested for the static verification of web service interface contracts [5, 8, 9].

LTL has been introduced to express properties about sequences of states in systems called Kripke structures [3]. In the current case, the states under consideration are XML objects termed messages. Let us denote by M the set of XML messages. A sequence of messages $m_1, m_2 \ldots$, where $m_i \in M$ for every $i \geq 1$, is called a message trace. We write m_i to denote the i-th message of the trace $\overline{m}$, and $\overline{m}^i$ to denote the trace obtained from $\overline{m}$ by starting at the i-th message.

A domain function is used to fetch and compare values inside a message; it receives an argument π representing a *path* from the root to some element of the message. This path is defined using standard, XPath 1.0 notation. Formally, if we let I be a domain of values, and Π be the set of XPath expressions, the domain function Dom is an application $M \times \Pi \rightarrow 2^I$ which, given a message $m \in M$ and a path $\pi \in \Pi$, returns a subset $Dom_m(\pi)$ of I, representing the set of values appearing in message m at the end of the path π. For example, if we let Π be the

set of XPath formulæ, $\pi \in \Pi$ be the particular formula "/message/stock/name", and $m \in M$ be the following message:

```
<message>
  <action>placeBuyOrder</action>
  <stock>
    <name>stock-1</name>
    <amount>123</amount>
  </stock>
  <stock>
    <name>stock-2</name>
    <amount>456</amount>
  </stock>
</message>
```

then $Dom_m(\pi) = \{\text{stock-1}, \text{stock-2}\}$.

LTL-FO$^+$'s syntax is based on classical propositional logic, using the connectives $\neg$ ("not"), $\vee$ ("or"), $\wedge$ ("and"), $\rightarrow$ ("implies"), to which four temporal operators have been added. An LTL-FO$^+$ formula is a well-formed combination of these operators and connectives, according to the usual construction rules:

Definition 1 (Syntax) *1. If x and y are variables or constants, then $x = y$ is a LTL-FO$^+$ formula;*
2. If φ and ψ are LTL-FO$^+$ formulæ, then $\neg\varphi$, $\varphi \wedge \psi$, $\varphi \vee \psi$, $\varphi \rightarrow \psi$, $\mathbf{G}\,\varphi$, $\mathbf{F}\,\varphi$, $\mathbf{X}\,\varphi$, $\varphi\,\mathbf{U}\,\psi$, $\varphi\,\mathbf{V}\,\psi$ are LTL-FO$^+$ formulæ;
3. If φ is a LTL-FO$^+$ formula, x_i is a free variable in φ, $p \in \Pi$ is a XPath value, then $\exists_p x_i : \varphi$ and $\forall_p x_i : \varphi$ are LTL-FO$^+$ formulæ.

The semantics of an LTL-FO$^+$ formula is given with respect to a partial function $p : V \rightarrow I$ that assigns every free variable in the formula. Let's denote by $P[p, \psi](\overline{m})$ the predicate that outputs **true** if the message trace $\overline{m}$ satisfies ψ given p and **false** otherwise. The semantics of LTL-FO$^+$ is then given as:

When p is clear from context we write $\overline{m} \models \varphi$ to indicate that the trace $\overline{m}$ satisfies φ. As usual, we define the semantics of the other connectors with the following identities: $\varphi \wedge \psi \equiv \neg(\neg\varphi \vee \neg\psi)$, $\varphi \rightarrow \psi \equiv \neg\varphi \vee \psi$, $\mathbf{G}\,\varphi \equiv \neg(\mathbf{F}\,\neg\varphi)$, $\varphi\,\mathbf{V}\,\psi \equiv \neg(\neg\varphi\,\mathbf{U}\,\neg\psi)$, $\forall_p x : \varphi \equiv \neg(\exists_p x : \neg\varphi)$.

Boolean connectives carry their usual meaning. The temporal operator **G** means "globally": the formula $\mathbf{G}\,\varphi$ means φ holds for every message of the trace. The operator **F** means "eventually"; the formula $\mathbf{F}\,\varphi$ holds if φ holds for some future message of the trace. The operator **X** means "next" and $\mathbf{X}\ \varphi$ holds whenever φ holds in the next message of the trace. Finally, the **U** operator means "until" and formula $\varphi\,\mathbf{U}\,\psi$ holds if φ holds for every messages until some message satisfies ψ.

Any LTL-FO$^+$ formula also has an equivalent *negation normal form*. An LTL-FO$^+$ formula is in *negation normal form* if it doesn't contain the operators $\mathbf{F}$, $\mathbf{G}$ and $\rightarrow$, and if all negations $\neg$ are pushed inside until they precede equalities. Definition : Negation normal form

$$
\begin{aligned}
P[p, c_1 = c_2](\overline{m}) &\Leftrightarrow p(c_1) = p(c_2)\\
P[p, c_1 \neq c_2](\overline{m}) &\Leftrightarrow p(c_1) \neq p(c_2)\\
P[p, \mu \vee \eta](\overline{m}) &\Leftrightarrow P[p, \mu](\overline{m}) \vee P[p, \eta](\overline{m});\\
P[p, \mu \wedge \eta](\overline{m}) &\Leftrightarrow P[p, \mu](\overline{m}) \wedge P[p, \eta](\overline{m});\\
P[p, \mathbf{X}\ \psi](\overline{m}) &\Leftrightarrow P[p, \psi](\overline{m}^2);\\
P[p, \mu\ \mathbf{U}\ \eta](\overline{m}) &\Leftrightarrow P[p, \eta](\overline{m}) \vee \Big(P[p, \mu](\overline{m}) \wedge P[p, \mu\ \mathbf{U}\ \eta](\overline{m}^2)\Big);\\
P[p, \mu\, \mathbf{R}\, \eta](\overline{m}) &\Leftrightarrow P[p, \mu \wedge \eta](\overline{m}) \vee \Big(P[p, \eta](\overline{m}) \wedge P[p, \mu\, \mathbf{R}\, \eta](\overline{m}^2)\Big);\\
P[p, \exists_\pi x : \psi](\overline{m}) &\Leftrightarrow \bigvee_{x_i \in Dom_m(\pi)} P[p \cup \{(x, x_i)\}, \psi](\overline{m});\\
P[p, \forall_\pi x : \psi](\overline{m}) &\Leftrightarrow \bigwedge_{x_i \in Dom_m(\pi)} P[p \cup \{(x, x_i)\}, \psi](\overline{m}).
\end{aligned}
$$

Table 1: Semantics of LTL-FO$^+$

Definition 2 *An LTL-FO$^+$ formula is in* negation normal form *if it doesn't contain the operators* $\mathbf{F}$, $\mathbf{G}$ *and* $\rightarrow$, *and if all negations* $\neg$ *are pushed inside until they precede equalities.*

We identify $\neg(x = y)$ with $x \neq y$ in order to eliminate the operator $\neg$ completely. This form is always obtainable with the use the identities $\mathbf{F}\ \psi \equiv \mathbf{true}\ \mathbf{U}\ \psi$, $\mathbf{G}\ \psi \equiv \mathbf{false}\, \mathbf{R}\, \psi$ and $\psi_1 \rightarrow \psi_2 \equiv \neg\psi_1 \vee \psi_2$, and by remembering that $\vee$, $\mathbf{U}$, and $\exists$ are the dual of $\wedge$, $\mathbf{R}$, and $\forall$ respectively. It what follows, we consider only properties in negation normal form.

The notion of *temporal depth* of formulæ serves as the basis for induction in a number of proofs.

Definition 3 *The* temporal depth *of an LTL-FO$^+$ formula φ, denoted depth(φ), is the maximal number of nodes associated with a temporal operator (*$\mathbf{F}$, $\mathbf{G}$, $\mathbf{X}$, $\mathbf{U}$ *or* $\mathbf{R}$*) that can be observed in a branch of φ.*

Variable assignments can be represented mathematically as a partial function from a set of variables to a set of values. We denote by V the set of variable that occur in a LTL-FO$^+$ formula and by I the domain of values that may appear in a message. We write $\mathcal{PF}_{V \rightarrow I}$, for the set of partial functions from V to I. The state of all variables in V during a run of the automaton A_φ can always be represented by a partial function $p : V \rightarrow I$. Some variables may not be assigned yet, but p is updated continuously as A_φ reads the input trace, with $p(x)$ representing the valuation of variable $x \in V$. Abusing the notation we write $p(v) = v$ for any constant value $v \in I$ that occurs in a formula.

2.1 Monitoring LTL-FO$^+$

By repeatedly applying the classical semantic rules of LTL, the evaluation of an LTL formula φ on a trace σ induces a tree. For example, in the case of the formula $\mathbf{G}\,(a \rightarrow \mathbf{X}\,b)$ evaluated on the trace cab, the top-level operator of that formula, $\mathbf{G}$, corresponds to the top-level node of the tree. According to the semantics of LTL, $\mathbf{G}\,\varphi$ is true if and only if φ is true for every suffix of the current trace. The tree hence spawns three child nodes, corresponding to the evaluation of $a \rightarrow Xb$ for traces cab, ab and b, respectively. Taking the first such child node, the the top-level operator now becomes $\rightarrow$; this operator evaluates to when, on the current trace, either a evaluates to $\bot$ or $\mathbf{X}\,b$ evaluates to $\top$. This, in turn, spawns two child nodes corresponding to each condition, and so on. Ultimately, only equalities on values remain, and the trough value of each subformula can then be obtained by combining and propagating values towards the top of the tree.

3 Related Representations

Büchi Automata [2], extend finite non-deterministic automata to infinite traces and provide an automata representation of LTL formulæ. Multiple translation algorithms exist [10], and the automaton's state set is linear in the size of the property under consideration. Unlike automata that recognize finite sequences, the nondeterminism of the automata is essential to its expressive power. An infinite sequence is valid if it enters an accepting state infinitely often on at least one of its possible runs.

Vardi [10], further suggested extending the Büchi Automata to include both existential and universal transitions. When encountering an existential choice for a given input token, the automata non-deterministically chooses one of them, and can then be though of as being in either one of its multiple possible destination states (thus behaving in the same manner as a non-deterministic Büchi automata). When encountering a universal transition, the automata's run continues simultaneously in each one of the universal transition' s destination state. A run of an alternating automata generates a tree, and an infinite input sequence is accepting of the *exists* at least one tree for which *every* branch visits an accepting state infinitely often. While alternating automata are no more expressive than regular Büchi automata, they can be exponentially more concise.

A more expressive representation was proposed by Barringer et al. [1]. The devised the Quantified Event Automata, an automata enriched with quantified variables. Each variable is associated with a domain of values that have been observed so far, used to determine acceptance. Compared with LTL-FO$^+$, QEA are strictly less expressive because QEA restrict the position on quantifiers in the formula.

Cassar et al. [4] introduce dynamic automata with timers and events (DATEs) which serves as basis for verification in the LARVA system. Using Dates, specific events can engender a duplication of the property automata under consideration.

For example, in a scenario in which multiple users interact with a service using a specific protocol, modeled by an automata. A new automaton instance would be generated each time a user logs onto the system and initiates the communication protocol. The expressivity of LTL-FO$^+$ is orthogonal to that of DATEs. The quantifiers of present in LTL-FO$^+$ formulæ allows it to process traces in which a single event contains multiple instances with the same name, which is not possible with DATEs. However, DATEs posses clocks and internal variables that allow them to verify some behaviors that cannot be stated with LTL-FO$^+$.

4 Automata-Theoretic Representation of LTL-FO$^+$

In this section, we show how, given an LTL-FO$^+$ formula φ, one can build a *modified alternating Büchi automaton* $A_\varphi = (\Sigma, V, I, S, s^0, \rho, F)$ such that the language recognized by A_φ is exactly the set of message traces satisfying φ. The alphabet Σ is the set M of all XML messages. The set V consists of all the variables that compose φ. The set C includes all the constants present φ as well as any value that can be assigned to the variables in V when the automaton reads a message $m \in \Sigma$. The set S of states consists of all subformulas of φ, denoted by $sub(\varphi)$, defined recursively as follows:

$\varphi \quad \in sub(\varphi)$;
$\mu \vee \eta \quad \in sub(\varphi) \Rightarrow \mu, \eta \in sub(\varphi)$;
$\mu \wedge \eta \quad \in sub(\varphi) \Rightarrow \mu, \eta \in sub(\varphi)$;
$\mathbf{X}\ \psi \quad \in sub(\varphi) \Rightarrow \ \psi \in sub(\varphi)$;
$\mu\ \mathbf{U}\ \eta \quad \in sub(\varphi) \Rightarrow \mu, \eta \in sub(\varphi)$;
$\mu\, \mathbf{R}\, \eta \quad \in sub(\varphi) \Rightarrow \mu, \eta \in sub(\varphi)$;
$\exists_\pi x : \psi \in sub(\varphi) \Rightarrow \ \psi \in sub(\varphi)$;
$\forall_\pi x : \psi \in sub(\varphi) \Rightarrow \ \psi \in sub(\varphi)$.

The set S additionally contains two distinct states, an accepting state $\top$ and its negation $\bot$. Both of these are "pit" states with any outgoing transition looping back to themselves.

A run of A_φ is characterized by the states of S that are visited as well as by the variable assignments that hold during these visits. Hence, the current "real" states of A_φ can be thought as a pair (p, ψ) from $\mathcal{PF}_{V \to I} \times \ S$ where p represents the variable assignments that currently holds at ψ. The transition function $\rho : \mathcal{PF}_{V \to I} \times \ S \times$ also operates on $\mathcal{PF}_{V \to I} \times \ S$ rather than on S.

The initial state $s^0 \in S$ is φ itself and the initial input of ρ is the couple $(\varnothing, s^0)$ where $\varnothing$ denotes the empty partial function from V to I. The set $F \subset S$ of accepting states includes $\top$ and every Release formula in S. The variable assignments do not affect the acceptance of the run: any couple (p, ψ) is accepting if $\psi \in F$. We write F_ψ to refer to the set of accepting states built from any formula ψ. More formally, if x and y are variables or constants, and if ψ, μ, and η are LTL-FO$^+$ formulæ, the set of accepting states is recursively defined according to the following rules:

- $F_{(x=y)} = F_{(x \neq y)} = \{\top\}$;
- $F_{(\mathbf{X}\ \psi)} = F_{(\exists_\pi x:\psi)} = F_{(\forall_\pi x:\psi)} = F_\psi$;

- $F_{(\mu \vee \eta)} = F_{(\mu \wedge \eta)} = F_{(\mu \ \mathbf{U} \ \eta)} = F_\mu \cup F_\eta$;
- $F_{(\mu \, \mathbf{R} \, \eta)} = F_\mu \cup F_\eta \cup \{\mu \, \mathbf{R} \, \eta\}$.

It remains only to define the transition function ρ. This is efficiently done by listing a small amount of general rules that must be applied recursively. Let p be a partial function from V to I, and m be any message in Σ. We set:

- $\rho\big((p, \top), m\big) = (\varnothing, \top)$;
- $\rho\big((p, \bot), m\big) = (\varnothing, \bot)$;
- $\rho\big((p, x = y), m\big) = \begin{cases} (\varnothing, \top) & \text{if } p(x) = p(y) \\ (\varnothing, \bot) & \text{otherwise;} \end{cases}$
- $\rho\big((p, x \neq y), m\big) = \begin{cases} (\varnothing, \top) & \text{if } p(x) \neq p(y) \\ (\varnothing, \bot) & \text{otherwise;} \end{cases}$
- $\rho\big((p, \mu \vee \eta), m\big) = \rho\big((p, \mu), m\big) \vee \rho\big((p, \eta), m\big)$;
- $\rho\big((p, \mu \wedge \eta), m\big) = \rho\big((p, \mu), m\big) \wedge \rho\big((p, \eta), m\big)$;
- $\rho\big((p, \mathbf{X} \ \psi), m\big) = (p, \psi)$;
- $\rho\big((p, \mu \ \mathbf{U} \ \eta), m\big) = \rho\big((p, \eta), m\big) \vee \Big(\rho\big((p, \mu), m\big) \wedge (p, \mu \ \mathbf{U} \ \eta)\Big)$;
- $\rho\big((p, \mu \, \mathbf{R} \, \eta), m\big) = \rho\big((p, \mu \wedge \eta), m\big) \vee \Big(\rho\big((p, \eta), m\big) \wedge (p, \mu \, \mathbf{R} \, \eta)\Big)$;
- $\rho\big((p, \exists_\pi x : \psi), m\big) = \bigvee_{x_i \in Dom_m(\pi)} \rho\big((p \cup \{(x, x_i)\}, \psi), m\big) \vee (\varnothing, \bot)$;
- $\rho\big((p, \forall_\pi x : \psi), m\big) = \bigwedge_{x_i \in Dom_m(\pi)} \rho\big((p \cup \{(x, x_i)\}, \psi), m\big) \wedge (\varnothing, \top)$.

If $Dom_m(\pi)$ is empty, we take the disjunction $\vee_{x_i \in Dom_m(\pi)}$ in the case $\exists$ to be equivalent to **false**, and the conjunction $\wedge_{x_i \in Dom_m(\pi)}$ in the case $\forall$ to be equivalent to **true**. Hence, the last two cases of the transition formula evaluate to $(\varnothing, \bot)$ and $(\varnothing, \top)$ respectively. These special rules are logically consistent with the inherent meaning of $\exists$ and $\forall$.

The constants **true** and **false**, which we include in the set I, may appear in the formulæ of S due to identities such as $\mathbf{F} \ \psi \equiv \mathbf{true} \ \mathbf{U} \ \psi$. In order for them to be compatible with the syntax of LTL-FO$^+$, and our definition of ρ, we identify **true** with the equality $\mathbf{true} = \mathbf{true}$, and **false** with the inequality $\mathbf{false} \neq \mathbf{false}$.

Note that ρ is undefined if its input contains an equality or inequality for which one of the variables are undefined in p. However, as long as φ, and by implication all of its subformulas are well-formed, this will never occur. Indeed, any variable in a well-formed LTL-FO$^+$ formula must be preceded by a quantifier on said variable. Therefore, ρ will first process the quantifier and assign values to this variable before reaching the equality or inequality.

Definition 4 *Let ψ be an LTL-FO$^+$ formula, and $p : V \to I$ be a partial function that assigns a value to every free variable in ψ. A* run of ρ on a message trace $\overline{m} = m_0, m_1, m_2, \ldots$ with root (p, ψ) *is an infinite $(\mathcal{PF}_{V \to I} \times S)$-labelled tree that respects two conditions:*

1. *The root node must is labelled (p, ψ);*
2. *Let's denote the distance between a node N and the root by $d(N)$, its label by $l(N)$, and the labels of its c_N children by $L_N = \{l_1, \ldots, l_{c_N}\}$. For any node N, the set L_N must satisfy $\rho\big(l(N), m_{d(N)}\big)$, and for any $1 \leq i \leq c_N$, the label (or couple) l_i must appear in $\rho\big(l(N), m_{d(N)}\big)$.*

For any LTL-FO$^+$ formula ψ, a run of ρ is said to be F_ψ-accepting *if and only if every branch of the run has an infinite number of nodes whose labels contain a state in F_ψ; in other words, if every branch* visits F_ψ infinitely often.

A "run of ρ" is a tree of couples from $\mathcal{PF}_{V\to D} \times S$ generated by following the rules of ρ for all messages in a trace. These runs are more general than the runs of a modified alternating Büchi automaton A_φ whose root is limited to the label $(\varnothing, \varphi)$. The previous definition allows us to properly name, and work with, parts of automaton runs, which happens frequently in Section 5. The parts that are themselves runs of ρ will be called *subruns*.

Definition 5 *Let ψ be any LTL-FO$^+$ formula. A* run of the automaton A_φ on a message trace $\overline{m}$ *is an run of ρ on $\overline{m}$ and with root $(\varnothing, \varphi)$. Such a run is* accepting *if and only if it is F_φ-accepting. A_φ* accepts $\overline{m}$ *if A_φ admits at least one accepting run on $\overline{m}$.*

5 Proof of Correctness

5.1 Preliminaries

The following lemma shows that any transition formula in ρ can be expressed as a disjunction of conjunctive clauses which themselves consist only in equalities, inequalities, couples in $\mathcal{PF}_{V\to I} \times S$ or couples of the form $(\varnothing, \top)$ or $(\varnothing, \bot)$. This form can be obtained once every recursive rule of ρ, except the ones on equalities and inequalities, has been applied and makes it easier to identify the set of states that may compose the next level of a run tree. In what follows, for a given disjunction D, the sets E_d ranges over the equalities and inequalities, N_d ranges over couples in $\mathcal{PF}_{V\to I} \times S$ and A_d ranges over couples of the form $(\varnothing, \top)$ or $(\varnothing, \bot)$. Intuitively, the disjunctive form represents the multiples combinations of states that are simultaneously visited by the automata during a run.

Lemma 1 *For any message $m \in M$, any LTL-FO$^+$ formula ψ, and any partial function $p : V \to I$ that assigns every free variable in ψ, $\rho\big((p, \psi), m\big)$ is equal to an expression of the form*

$$\bigvee_{d\in D} \Big(\bigwedge_{e\in E_d} \rho\big((p_{d,e}, \varepsilon_{d,e}), m\big) \wedge \bigwedge_{n\in N_d} (p_{d,n}, \psi_{d,n}) \wedge \bigwedge_{a\in A_d} (\varnothing, \top_{d,a})\Big). \tag{1}$$

For all $d \in D$, $e \in E_d$, $n \in N_d$, and $a \in A_d$, $\varepsilon_{d,e}$ is an equality or inequality, $\psi_{d,n}$ is a subformula of ψ (which includes ψ), $\top_{d,a}$ is either $\top$ or $\bot$, and $p_{d,e}$ and $p_{d,n}$ are partial functions from V to I that assign every free variable in $\varepsilon_{d,e}$ and $\psi_{d,n}$ respectively.

Moreover, if $depth(\psi) = 0$, then N_d is empty for all $d \in D$.Otherwise, N_d may not always be empty, and any $\psi_{d,n}$ fits only one of three descriptions:

1) $depth(\psi_{d,n}) < depth(\psi)$;
2) $depth(\psi_{d,n}) = depth(\psi)$ and $\psi_{d,n} = \mu \mathbf{U} \eta$ for some $\mu, \eta \in sub(\psi)$;

3) $depth(\psi_{d,n}) = depth(\psi)$ and $\psi_{d,n} = \mu\,\mathbf{R}\,\eta$ for some μ, $\eta \in sub(\psi)$.

Observe that, since the rules governing the decomposition of equalities and inequalities are not appliqued when decomposing a formula in Normal form, if $depth(\psi) = 0$, then N_d and A_d is empty for all $d \in D$. Indeed, elements are only be added in a set A_d in the particular case when an a qualifier $\forall$ (resp. $\exists$) is encountered with an empty domain, leading to a trivial true (resp. false) verdict.

Remark 1 *For the sake of conciseness, we will regularly shorten the decomposition described in Lemma 1 to:*

$$\rho\big((p,\psi),m\big) = \bigvee_{d\in D}\Big(\bigwedge_{c\in C_d} R_{d,c}\Big).$$

The set C_d, for all $d \in D$, includes every index in E_d, N_d, and A_d. The terms $R_{d,c}$, are of the form $\rho\big((p_{d,e},\varepsilon_{d,e}),m\big)$, $(p_{d,n},\psi_{d,n})$, or $(\varnothing,\top_{d,a})$.

Proof (Proof of Lemma 1). Each term $R_{d,c}$ in the decomposition of $\rho\big((p,\psi),m\big)$ is obtained by recursively applying the rules of ρ. Every rule, if applicable, must be used exhaustively, except the rules on equalities and inequalities, which are not applied at this point. This restriction guarantees that terms of the form $\rho\big((p_{d,e},\varepsilon_{d,e}),m\big)$ are preserved. Thus, terms of the form $(\varnothing,\top)$ or $(\varnothing,\bot)$ only occur after evaluating the the quantifiers $\exists_\pi$ and $\forall_\pi$ when $Dom_m(\pi) = \varnothing$.

It follows from the restriction above, and the rules of ρ, that $\rho\big((p,\psi),m\big)$ can be recursively decomposed into a set consisting of couples $(f,\phi) \in \mathcal{PF}_{V\to I} \times S$ as well as of terms of the form $\rho\big((p_{d,e},\varepsilon_{d,e}),m\big)$. Any output of ρ that does not match one of these forms may be further decomposed. However, the recursive decomposition process cannot go on indefinitely, because while the rules of ρ sometimes increase the number of terms, the output states always have a smaller depth.

Moreover, a couple of the form (f,ϕ) only appears if ψ contains temporal operators or quantifiers. If (f,ϕ) is output by a rule related to the operator $\mathbf{X}$, $\mathbf{U}$ or $\mathbf{R}$, then the state ϕ is a subformula of ψ. Otherwise, if a couple (f,ϕ) is generated to a quantifier $\exists_\pi$ or $\forall_\pi$, $f = \varnothing$ and $\phi \in \{\top,\bot\}$. Hence, (f,ϕ) is either of the form $(p_{d,n},\psi_{d,n})$, $(\varnothing,\top_{d,a})$ or $(\varnothing,\bot_{d,a})$ and can thus be indexed by a set E_d, N_d, or A_d. It remains to show that these terms can be arranged into a disjunction of conjunctions.

We begin by showing that $\rho\big((p,\psi),m\big)$, and all of its partial decompositions, match a form that is similar, but not identical, to the desired one. It is still a disjunction of conjunctions, but its terms are partitioned by sets T_d and $C_d = E_d \cup A_d \cup N_d$ instead of just C_d. More precisely, we consider the form

$$\bigvee_{d\in D}\Big(\bigwedge_{t\in T_d}\rho\big((p_{d,t},\psi_{d,t}),m\big) \wedge \bigwedge_{c\in C_d} R_{d,c}\Big). \tag{2}$$

For all $d \in D$, and $t \in T_d$, $\psi_{d,t}$ is a subformula of ψ that is neither an equality nor inequality, and $p_{d,t}$ is a partial function from V to I that assigns every free

variable in $\psi_{d,t}$. In short, the sets T_d index terms that may be further decomposed. The terms $R_{d,c}$, which are either of the form $\rho\big((p_{d,e}, \varepsilon_{d,e}), m\big)$, $(p_{d,n}, \psi_{d,n})$, $(\varnothing, \top_{d,a})$ or $(\varnothing, \bot_{d,a})$, already satisfy the conditions of the lemma. In particular, the partial function of any term $R_{d,c}$ covers every free variable in its associated formula.

We prove the validity of this claim by induction. First, let's consider the initial term, $\rho\big((p, \psi), m\big)$, for which p assigns every free variable in ψ. If ψ is an equality or inequality, then this term is of the form $\rho\big((p_{d,e}, \psi_{d,e}), m\big)$. Otherwise, it matches the form $\rho\big((p_{d,t}, \psi_{d,t}), m\big)$. In both cases, $\rho\big((p, \psi), m\big)$ is the sole operand of a conjunction indexed by a set T_d or C_d where $D = \{d\}$. Therefore, the base case holds.

For the induction step, we show that the partial decomposition of $\rho\big((p, \psi), m\big)$, possibly $\rho\big((p, \psi), m\big)$ itself also is of the form (2). If $T_d = \varnothing$ for every $d \in D$, then the decomposition only contains terms of the form $R_{d,c}$. Hence, it matches the expression claimed by this lemma, as well as the form (2). Otherwise, if $T_{d^*} \neq \varnothing$ for some $d^* \in D$, then it contains a term of the form $\rho\big((p_{d^*,t^*}, \psi_{d^*,t^*}), m\big)$ that is indexed by d^* and some $t^* \in T_{d^*}$. Since $\rho\big((p_{d^*,t^*}, \psi_{d^*,t^*}), m\big)$ is further decomposable, we apply the corresponding rule of ρ to it. We prove below that the resulting terms can always be arranged in a way that preserves the disjunctive form (2). Therefore, the induction step also holds.

The details of maintaining the disjunctive form depend on the nature of the formula $\psi_{d,^*,t^*}$. If it is a Next formula $\mathbf{X}\ \mu$, then $\rho\big((p_{d^*,t^*}, \psi_{d^*,t^*}), m\big)$ outputs a couple (p_{d^*,t^*}, μ) that is simply added to the conjunction indexed by $N_{d^*} \subseteq C_{d^*}$. If ψ_{d^*,t^*} is of the form $\exists_\pi x : \mu$ or $\forall_\pi x : \mu$, but $Dom_m(\pi) = \varnothing$, then $\rho\big((p_{d^*,t^*}, \psi_{d^*,t^*}), m\big)$ outputs $(\varnothing, \bot)$ or $(\varnothing, \top)$ respectively. This output is then added to the conjunction indexed by $A_{d^*} \subseteq C_{d^*}$. The remaining cases, which concern the forms $\mu \vee \eta$, $\mu \wedge \eta$, $\mu\ \mathbf{U}\ \eta$, and $\mu\, \mathbf{R}\, \eta$, but also $\exists_\pi x : \mu$ and $\forall_\pi x : \mu$ when $Dom_m(\pi) \neq \varnothing$, are covered below.

Case $\psi_{d^*,t^*} = \forall_\pi x : \mu$ or $\psi_{d^*,t^*} = \mu \wedge \eta$ ($\mu, \eta \in sub(\psi)$, $Dom_m(\pi) \neq \varnothing$)

Let's first suppose that $\psi_{d^*,t^*} = \forall_\pi x : \mu$. By applying the corresponding rule of ρ on $\rho\big((p_{d^*,t^*}, \psi_{d^*,t^*}), m\big)$, and by considering the terms surrounding its output, we obtain the following equation:

$$\Big(\rho\big((p_{d^*,t^*}, \forall_\pi x : \mu), m\big) \wedge \bigwedge_{\substack{t \in T_{d^*} \\ t \neq t^*}} \rho\big((p_{d^*,t}, \psi_{d^*,t}), m\big) \wedge \bigwedge_{c \in C_{d^*}} R_{d^*,c}\Big) \vee \bigvee_{\substack{d \in D \\ d \neq d^*}} \dots$$

$$= \Big(\bigwedge_{x_i \in Dom_m(\pi)} \rho\big((p_{d^*,t^*} \cup \{(x, x_i)\}, \mu), m\big) \wedge \bigwedge_{\substack{t \in T_{d^*} \\ t \neq t^*}} \rho\big((p_{d^*,t}, \psi_{d^*,t}), m\big) \wedge \bigwedge_{c \in C_{d^*}} R_{d^*,c}\Big) \vee \bigvee_{\substack{d \in D \\ d \neq d^*}} \dots$$

The term $\rho\big((p_{d^*,t^*}, \psi_{d^*,t^*}), m\big)$, which belongs in the conjunction indexed by T_{d^*}, is replaced by a conjunction of $|Dom_m(\pi)|$ terms. Every one of them can be indexed by either T_{d^*} or E_{d^*} depending on whether μ is an equality, an inequality, or neither. This is because the partial function $p_{d^*,t^*} \cup \{(x, x_i)\}$, for any value

x_i, assigns every free variable in μ. The formulas μ and $\forall_\pi x : \mu$ share the same variables, with the only difference being that x is not bound inside μ. As such, the free variables of μ consists of x and the free variables of ψ_{d^*,t^*}, which is exactly what $p_{d^*,t^*} \cup \{(x, x_i)\}$ covers. Hence, the right-hand expression of the previous equation matches form (2).

The subcase $\psi_{d^*,t^*} = \mu \wedge \eta$, in which we decompose $\rho\big((p_{d^*,t^*}, \psi_{d^*,t^*}), m\big)$ into a conjunction of two terms rather than $|Dom_m(\pi)|$, is treated similarly. Obviously, the partial function p_{d^*,t^*} assigns every free variable in μ and η.

Case $\psi_{d^*,t^*} = \exists_\pi x : \mu$ or $\psi_{d^*,t^*} = \mu \vee \eta$ $(\mu, \eta \in sub(\psi),\ Dom_m(\pi) \neq \varnothing)$

Let's first suppose that $\psi_{d^*,t^*} = \exists_\pi x : \mu$. As with the subcase $\forall_\pi x : \mu$, ρ decomposes the term $\rho\big((p_{d^*,t^*}, \psi_{d^*,t^*}), m\big)$ into $|Dom_m(\pi)|$ terms of the form $\rho\big((p_{d^*,t^*} \cup \{(x, x_i)\}, \mu), m\big)$. However, unlike the subcase $\forall_\pi x : \mu$, these terms are joined by the operator $\vee$. If $|Dom_m(\pi)| \geq 2$, they create a disjunction inside the conjunction $d^* \in D$. As shown by the second member of the equation below, the resulting expression is not a conjunction of disjunctions.

$$\Big(\rho\big((p_{d^*,t^*}, \exists_\pi x : \mu), m\big) \wedge \bigwedge_{\substack{t \in T_{d^*} \\ t \neq t^*}} \rho\big((p_{d^*,t}, \psi_{d^*,t}), m\big) \wedge \bigwedge_{c \in C_{d^*}} R_{d^*,c}\Big) \vee \bigvee_{\substack{d \in D \\ d \neq d^*}} \dots$$

$$= \Big(\bigvee_{x_i \in Dom_m(\pi)} \rho\big((p_{d^*,t^*} \cup \{(x, x_i)\}, \mu), m\big) \wedge \bigwedge_{\substack{t \in T_{d^*} \\ t \neq t^*}} \rho\big((p_{d^*,t}, \psi_{d^*,t}), m\big) \wedge \bigwedge_{c \in C_{d^*}} R_{d^*,c}\Big) \vee \bigvee_{\substack{d \in D \\ d \neq d^*}} \dots$$

$$= \bigvee_{x_i \in Dom_m(\pi)} \Big(\rho\big((p_{d^*,t^*} \cup \{(x, x_i)\}, \mu), m\big) \wedge \bigwedge_{\substack{t \in T_{d^*} \\ t \neq t^*}} \rho\big((p_{d^*,t}, \psi_{d^*,t}), m\big) \wedge \bigwedge_{c \in C_{d^*}} R_{d^*,c}\Big) \vee \bigvee_{\substack{d \in D \\ d \neq d^*}} \dots$$

Therefore, we distribute the surrounding conjunctions, which are indexed by the sets $T_{d^*} \setminus \{t^*\}$ and C_{d^*}, over the problematic disjunction. This results in $|Dom_m(\pi)|$ almost identical conjunctions that differ by the value x_i inside their term $\rho\big((p_{d^*,t^*} \cup \{(x, x_i)\}, \mu), m\big)$. As shown by the last member of the equation above, this new expression is a disjunction of conjunctions.

We established, with the subcase $\forall_\pi x : \mu$, that the partial function $p_{d^*,t^*} \cup \{(x, x_i)\}$, for any value x_i, assigns every free variable in μ. Thus, every term of the form $\rho\big((p_{d^*,t^*} \cup \{(x, x_i)\}, \mu), m\big)$, depending on the nature of μ, can be indexed by the set T_{d^*} or $E_{d^*} \subseteq C_{d^*}$ associated with its conjunction. Hence, the expression obtained matches form (2).

The subcase $\psi_{d^*,t^*} = \mu \vee \eta$, in which we decompose $\rho\big((p_{d^*,t^*}, \psi_{d^*,t^*}), m\big)$ into a disjunction of two terms rather than $|Dom_m(\pi)|$, is treated similarly.

Case $\psi_{d^*,t^*} = \mu\ \mathbf{U}\ \eta$ or $\psi_{d^*,t^*} = \mu\, \mathbf{R}\, \eta$ $(\mu, \eta \in sub(\psi))$

Let's first consider the case where $\psi_{d^*,t^*} = \mu\ \mathbf{U}\ \eta$. By the definition of ρ $\rho\big((p_{d^*,t^*}, \psi_{d^*,t^*}), m\big)$ is decomposable into a disjunction of two terms, namely

$\rho\big((p_{d^*,t^*},\eta),m\big)$ and $\rho\big((p_{d^*,t^*},\mu),m\big) \wedge (p_{d^*,t^*},\mu\ \mathbf{U}\ \eta)$. It follows from the previous case that the conjunctions surrounding them must be distributed over the disjunction. The process is displayed here:

$$\Bigg(\rho\big((p_{d^*,t^*},\mu\ \mathbf{U}\ \eta),m\big) \wedge \bigwedge_{\substack{t\in T_{d^*}\\ t\neq t^*}} \rho\big((p_{d^*,t},\psi_{d^*,t}),m\big) \wedge \bigwedge_{c\in C_{d^*}} R_{d^*,c}\Bigg) \vee \bigvee_{\substack{d\in D\\ d\neq d^*}} \dots$$

$$= \Bigg(\rho\big((p_{d^*,t^*},\eta),m\big) \vee \Big(\rho\big((p_{d^*,t^*},\mu),m\big) \wedge (p_{d^*,t^*},\mu\ \mathbf{U}\ \eta)\Big) \wedge \bigwedge_{\substack{t\in T_{d^*}\\ t\neq t^*}} \rho\big((p_{d^*,t},\psi_{d^*,t}),m\big) \wedge \bigwedge_{c\in C_{d^*}} R_{d^*,c}\Bigg) \vee \bigvee_{\substack{d\in D\\ d\neq d^*}} \dots$$

$$= \Bigg(\rho\big((p_{d^*,t^*},\eta),m\big) \wedge \bigwedge_{\substack{t\in T_{d^*}\\ t\neq t^*}} \rho\big((p_{d^*,t},\psi_{d^*,t}),m\big) \wedge \bigwedge_{c\in C_{d^*}} R_{d^*,c}\Bigg) \vee \Bigg(\rho\big((p_{d^*,t^*},\mu),m\big) \wedge (p_{d^*,t^*},\mu\ \mathbf{U}\ \eta) \wedge \bigwedge_{\substack{t\in T_{d^*}\\ t\neq t^*}} \rho\big((p_{d^*,t},\psi_{d^*,t}),m\big) \wedge \bigwedge_{c\in C_{d^*}} R_{d^*,c}\Bigg) \vee \bigvee_{\substack{d\in D\\ d\neq d^*}} \dots$$

The partial function p_{d^*,t^*} assigns every free variable inside the subformulas μ and η. Thus, the output terms $\rho\big((p_{d^*,t^*},\mu),m\big)$ and $\rho\big((p_{d^*,t^*},\eta),m\big)$, depending on the nature of μ and η, can be indexed by their respective copy of T_{d^*} or $E_{d^*} \subseteq C_{d^*}$. The couple $(p_{d^*,t^*},\mu\ \mathbf{U}\ \eta)$, by comparison, can be indexed by of $N_{d^*} \subseteq C_{d^*}$. Thus, the resulting expression is of the form (2).

The subcase $\psi_{d^*,t^*} = \mu\,\mathbf{R}\,\eta$ is treated similarly since the decomposition of $\rho\big((p_{d^*,t^*},\mu\,\mathbf{R}\,\eta),m\big)$ is nearly identical. Since p_{d^*,t^*} assigns every free variable in $\mu\,\mathbf{R}\,\eta$, it does the same for η and $\mu \wedge \eta$. Note though that while μ and η are subformulas of ψ because $\mu\,\mathbf{R}\,\eta \in sub(\psi)$, this is usually not the case for $\mu \wedge \eta$. Therefore, the term $\rho\big((p_{d^*,t^*},\mu \wedge \eta),m\big)$ cannot usually be indexed by the set T_{d^*}, but its decomposition $\rho\big((p_{d^*,t^*},\mu),m\big) \wedge \rho\big((p_{d^*,t^*},\eta),m\big)$ always can.

The above reasoning proves that any intermediate decomposition of the term $\rho\big((p,\psi),m\big)$ (with the rules for (in)equality omitted) will be of the form (2). However, we also showed that the induction step cannot be repeated indefinitely, since the depth of the terms monotonically decreases and the set T_d eventually becomes empty for every conjunction $d \in D$. When this happens, the resulting equation is a conjunction of disjunction of the form described in described in Lemma 1. □

A very similar decomposition can be applied to a predicate representing an LTL-FO$^+$ formula ψ. As expected, this notation follows the semantics of LTL-FO$^+$. By convention, if $Dom_m(\pi)$ is empty, then disjunctions and conjunctions indexed by $x_i \in Dom_m(\pi)$ default to **false** and **true** respectively. We keep

this convention, but we also propose a cosmetic change that allows the above definition to better match the definition of ρ.

$P[p, \top](\overline{m}) = \textbf{true}$ and $P[p, \bot](\overline{m}) = \textbf{false}$ for any partial function $p : V \to I$ and any trace $\overline{m}$. Hence, if $Dom_m(\pi)$ is empty, we have that $P[p, \exists_\pi x : \psi](\overline{m})$ and $P[p, \forall_\pi x : \psi](\overline{m})$ are equivalent to $P[\varnothing, \bot](\overline{m}^2)$ and $P[\varnothing, \top](\overline{m}^2)$ respectively. This creates an equivalences between the predicate rules and the rules of ρ. The decomposition of a predicate for a formula ψ is given in lemma 2.

Lemma 2 *For any message trace $\overline{m} = m_1, m_2, \ldots$, any LTL-FO$^+$ formula ψ, and any partial function $p : V \to I$ that assigns every free variable in ψ, $P[p, \psi](\overline{m})$ is equivalent to an expression of the form*

$$\bigvee_{d \in D} \Big(\bigwedge_{e \in E_d} P[p_{d,e}, \varepsilon_{d,e}](\overline{m}) \wedge \bigwedge_{n \in N_d} P[p_{d,n}, \psi_{d,n}](\overline{m}^2) i \wedge \bigwedge_{a \in A_d} P[\varnothing, \top_{d,a}](\overline{m}^2) \Big). \quad (3)$$

For all $d \in D$, $e \in E_d$, $n \in N_d$, and $a \in A_d$, $\varepsilon_{d,e}$ is an equality or inequality, $\psi_{d,n}$ is a subformula of ψ (which includes ψ), $\top_{d,a}$ is either $\top$ or $\bot$, and $p_{d,e}$ and $p_{d,n}$ are partial functions from V to I that assign every free variable in $\varepsilon_{d,e}$ and $\psi_{d,n}$ respectively.

Moreover, if $depth(\psi) = 0$, then N_d is empty for all $d \in D$. Otherwise, N_d may not always be empty, and any $\psi_{d,n}$ fit only one of three descriptions:

1) $depth(\psi_{d,n}) < depth(\psi)$;
2) $depth(\psi_{d,n}) = depth(\psi)$ and $\psi_{d,n} = \mu \, \mathbf{U} \, \eta$ for some μ, $\eta \in sub(\psi)$;
3) $depth(\psi_{d,n}) = depth(\psi)$ and $\psi_{d,n} = \mu \, \mathbf{R} \, \eta$ for some μ, $\eta \in sub(\psi)$.

Proof. Completely analogous to the proof of Lemma 1. □

Lemmas 1 and 2 identify a very strong connection between the semantics of LTL-FO$^+$ and the transition function ρ. In fact, it should come as no surprise, at this point, that their decompositions can be paired in such a way that they match, term for term.

Lemma 3 *For any message trace $\overline{m} = m, m_2, m_3, \ldots$, any LTL-FO$^+$ formula ψ, and any partial function $p : V \to I$ that assigns every free variable in ψ, the terms $p_{d,e}$, $p_{d,n}$, $\varepsilon_{d,e}$, $\psi_{d,n}$, and $\top_{d,a}$ in the decompositions of $\rho\big((p, \psi), m\big)$ and $P[p, \psi](\overline{m})$ can be labelled in such a way that they perfectly match across the two decompositions for any $d \in D$, $e \in E_d$, $n \in N_d$, and $a \in A_d$.*

5.2 A_φ accepts $\overline{m} \Rightarrow \overline{m} \models \varphi$

The proof relies upon the following theorem, which is straightforwardly derived from the definitions of $\mathbf{U}$ and $\mathbf{R}$ and describes the shape of the tree resulting from a run of $\rho\big((p, \psi), \overline{m}\big)$ with $\psi = \mathbf{U}$ or $\psi = \mathbf{R}$.

Lemma 4 *Let ψ, μ, and η be LTL-FO$^+$ formulæand $p : V \to I$ be a partial function that assigns every free variable in ψ. Suppose there exists an F_ψ-accepting run of ρ on a trace $\overline{m}$ with root (p, ψ). Let's denote this run by $r\big((p, \psi), \overline{m}\big)$.*

1) If $\psi = \mu\ \mathbf{U}\ \eta$, then:
- $\exists\ j \geq 1$ *: ρ admits an F_η-accepting run $r\big((p,\eta),\overline{m}^j\big)$ with root (p,η);*
- $\forall\ 1 \leq i < j$, *ρ admits an F_μ-accepting run $r\big((p,\mu),\overline{m}^i\big)$ with root (p,μ).*

2) If $\psi = \mu\,\mathbf{R}\,\eta$, then either:
- $\forall\ i \geq 1$, *ρ admits an F_η-accepting run $r\big((p,\eta),\overline{m}^i\big)$ with root (p,η); or*
- $\exists\ j \geq 1$ *: ρ admits an $F_{\mu\wedge\eta}$-accepting run $r\big((p,\mu\wedge\eta),\overline{m}^j\big)$ with root $(p,\mu\wedge\eta)$;*
- $\forall\ 1 \leq i < j$, *ρ admits an F_η-accepting run $r\big((p,\eta),\overline{m}^i\big)$ on $\overline{m}^i$ with root (p,η).*

Informally, the above lemma states that, if the automata admits an accepting run over $\phi = \mu\ \mathbf{U}\ \eta$ (resp. $\phi = \mu\,\mathbf{R}\,\eta$) then it admits accepting runs for μ and η on such prefixes and/or suffixes of the trace $\overline{m}$ as to be consistent with the semantics of these operators.

Proof. Because $\mathbf{U}$ and $\mathbf{R}$ lead to similar proofs, we give the details for $\mathbf{U}$, and focus on the key differences that occur with $\mathbf{R}$.

1) $\psi = \mu\ \mathbf{U}\ \eta$

The first step of the derivation of run $r\big((p,\psi),\overline{m}\big)$ is dictated by the rule

$$\rho\big((p,\mu\ \mathbf{U}\ \eta),m\big) = \rho\big((p,\eta),m\big) \vee \Big(\rho\big((p,\mu),m\big) \wedge (p,\mu\ \mathbf{U}\ \eta)\Big).$$

It is clear from the premise that p assigns every free variable in μ and η. Thus, we can apply Lemma 1 to both $\rho\big((p,\eta),m\big)$ and $\rho\big((p,\mu),m\big)$, which yields:

$$\rho\big((p,\mu\ \mathbf{U}\ \eta),m\big) = \bigvee_{d\in D_\eta}\Big(\bigwedge_{c\in C_d} R_{d,c}\Big) \vee \bigvee_{d\in D_\mu}\Big(\bigwedge_{c\in C_d} R_{d,c} \wedge (p,\mu\ \mathbf{U}\ \eta)\Big). \quad (4)$$

The sets D_η and D_μ encompass the decomposition of $\rho\big((p,\eta),m\big)$ and $\rho\big((p,\mu),m\big)$ respectively.

By definition 4, in a run of ρ, the children of a node $(p,\mu\ \mathbf{U}\ \eta)$ must include every couple $R_{d,c} \in D_\eta \cup D_\mu$ in the right-hand side of (4) Therefore, only two scenarios are possible:

(i) For any $i \geq 1$, there exists a distinguished index d_i such that message m_i is in D_μ;
(ii) For some $j \geq 1$, there exists a distinguished index d_j such that message m_j is in D_η, and for any $1 \leq i < j$, an index d_i such that message m_i is in D_μ.

The first scenario would occur if μ is verified by every message in $\overline{m}$ while the second occurs if μ holds continuously until some point j, where η holds (these scenarios are not mutually exclusive because two or more indexes can be represented by the child nodes, but cover every possibility). Indeed, either some $j \geq 1$ linked to D_η exists, in which case, the smallest possible j satisfies (ii); or it does not. In that case, (i) holds. Note also that every time a chosen index is

in D_μ, a node $(p, \mu \mathbf{U} \eta)$ appears as a child, so D_μ or D_η must be picked for the next message as well.

It turns out that the run $r\big((p,\psi),\overline{m}\big)$ follows scenario (ii). If it did not, then according to (i), and the previous discussion, there would be a branch in $r\big((p,\psi),\overline{m}\big)$ whose nodes are always $(p, \mu \mathbf{U} \eta)$. Since $\psi = \mu \mathbf{U} \eta$ is not in F_ψ, the run would not be F_ψ-accepting, thus contradicting our premise. Scenario (ii) echoes the desired results in regard to values i, j and formulas μ, η. It remains to extract, for any relevant value i or j, the desired subrun from $r\big((p,\psi),\overline{m}\big)$. We only cover the details for i, as both cases lead to almost identical proofs.

Since index d_i is in D_μ, the children of a node $(p, \mu \mathbf{U} \eta)$ linked to message m_i can be of three "types". They can be either $(p, \mu \mathbf{U} \eta)$ (type 1, mandatory), a couple $R_{d,c}$ in equation (4) where $d = d_i$ (type 2, mandatory), or a couple $R_{d,c}$ in (4) where $d \neq d_i$ (type 3, optional). To make sense of type 3, remember that an accepting run may include extra nodes as long as it respects ρ and F. Now, observe that the union of all children of type 2, denoted U_2^i, satisfies the conjunction $\wedge_{c \in C_{d_i}}$ in (4) if $m = m_i$. Thus, U_2^i satisfies $\rho\big((p,\mu), m_i\big)$. Hence, if we take a root node (p,μ), set U_2^i as its children, and keep every subrun of $r\big((p,\psi),\overline{m}\big)$ whose root is in U_2^i, we get a run $r\big((p,\mu),\overline{m}^i\big)$ starting at m_i and with root (p,μ). We must now prove that $r\big((p,\mu),\overline{m}^i\big)$ is F_μ-accepting. Since the original run is $F_{(\mu \mathbf{U} \eta)}$-accepting, all of its branches visit $F_{(\mu \mathbf{U} \eta)}$ infinitely often. This remains true for every branch in $r\big((p,\mu),\overline{m}^i\big)$ because they are infinite suffixes of branches in the original run (i is finite). However, the formulas in $r\big((p,\mu),\overline{m}^i\big)$ are limited to subformulas of μ (Lemma 1), so what is visited infinitely often is actually $F_{(\mu \mathbf{U} \eta)} \cap (sub(\mu) \cup \{\top\}) = F_\mu$.

As we mentioned earlier, the proof for j is very similar. The set U_2^j, where $d_j \in D_\eta$, identifies the desired subrun. The only notable difference is that a node $(p, \mu \mathbf{U} \eta)$, for message m_j, is not required to have a child of type 1.

2) $\psi = \mu \mathbf{R} \eta$

The beginning of the run $r\big((p,\psi),\overline{m}\big)$ is dictated by the rule

$$\rho\big((p, \mu \mathbf{R} \eta), m\big) = \rho\big((p, \mu \wedge \eta), m\big) \vee \Big(\rho\big((p,\eta),m\big) \wedge (p, \mu \mathbf{R} \eta)\Big),$$

which also applies each time the couple $(p, \mu \mathbf{R} \eta)$ appears in the run. As with $\mathbf{U}$, we can apply Lemma 1 and distribute $(p, \mu \mathbf{R} \eta)$ over the decomposition of $\rho\big((p,\eta),m\big)$ to get

$$\rho\big((p, \mu \mathbf{R} \eta), m\big) = \bigvee_{d \in D_{(\mu \wedge \eta)}} \Big(\bigwedge_{c \in C_d} R_{d,c} \Big) \vee \bigvee_{d \in D_\eta} \Big(\bigwedge_{c \in C_d} R_{d,c} \wedge (p, \mu \mathbf{R} \eta) \Big) \quad (5)$$

where the set $D_{(\mu \wedge \eta)}$ encompass the decomposition of $\rho\big((p, \mu \wedge \eta), m\big)$.

The scenarios (i) and (ii) introduced for $\mu \mathbf{U} \eta$ also apply in this case, but with minor adaptations regarding the index sets:

(i) For any $i \geq 1$, there exists a distinguished index d_i such that message m_i is in D_η;

(ii) For some $j \geq 1$, there exists a distinguished index d_j such that message m_j is in $D_{(\mu \wedge \eta)}$, and for any $1 \leq i < j$, there exists a distinguished index d_i such that message m_i is in D_η.

However, by comparison to $\mu \,\mathbf{U}\, \eta$, it is possible for $r\big((p,\psi),\overline{m}\big)$ to follow scenario (i). This is because ψ, this time, is included in F_ψ due to being a Release formula. Until formulas, by comparison, are not accepting. Thus, even though (i) implies the existence of a branch in $r\big((p,\psi),\overline{m}\big)$ whose nodes are always $(p, \mu \,\mathbf{R}\, \eta)$, such a branch visits F_ψ infinitely often and does not contradict our premise that the run is F_ψ-accepting. In fact, (i) and (ii) together echo the result stated in this lemma. The remaining of the proof for i is analogous to the one for i with $\mathbf{U}$. The same goes for j.

While every branch of the derived runs $r\big((p, \mu \wedge \eta), \overline{m}^j\big)$ and $r\big((p,\eta),\overline{m}^i\big)$ visits F_ψ infinitely often, their formulas are restricted to subformulas of $\mu \wedge \eta$ and η respectively (Lemma 1). Thus, the sets that are visited infinitely often are actually $F_\psi \cap (sub(\mu \wedge \eta) \cup \{\top\}) = F_{(\mu \wedge \eta)}$ and $F_\psi \cap (sub(\eta) \cup \{\top\}) = F_\eta$. □

We can now state the main lemmas of correction and completeness, indicating that an automata A_φ admits an accepting run for sequence ψ iff $\overline{m} \models \psi$.

Lemma 5 *For any LTL-FO$^+$ formula ψ and any partial function $p : V \to I$ that assigns every free variable in ψ, if there exists an F_ψ-accepting run of ρ on a trace $\overline{m} = m_1, m_2, \ldots$ and with root (p,ψ), then $P[p,\psi](\overline{m})$ is true.*

Proof. We proceed by strong induction on the temporal depth of LTL-FO$^+$ formulæ. Recall that by Lemmas 1 and 2, the decomposition of $\rho\big((p,\psi),m_1\big)$ and $P[p,\psi](\overline{m})$ can be given in normal disjunctive form as follows and that Lemma 3 assures us that for every $d \in D$, the terms of these two formulas can be matched to one another.

$$\rho\big((p,\psi),m_1\big) = \bigvee_{d \in D} \Big(\bigwedge_{e \in E_d} \rho\big((p_{d,e},\varepsilon_{d,e}),m_1\big) \wedge \bigwedge_{n \in N_d} (p_{d,n},\psi_{d,n}) \wedge \bigwedge_{a \in A_d} (\varnothing, \top_{d,a}) \Big) \tag{6}$$

$P[p,\psi](\overline{m}) \Leftrightarrow$

$$\bigvee_{d \in D} \Big(\bigwedge_{e \in E_d} P[p_{d,e},\varepsilon_{d,e}](\overline{m}) \wedge \bigwedge_{n \in N_d} P[p_{d,n},\psi_{d,n}](\overline{m}^2) \wedge \bigwedge_{a \in A_d} P[\varnothing,\top_{d,a}](\overline{m}^2) \Big) \tag{7}$$

Base case: Lemma 5 holds for any formula of depth 0.

Let $depth(\psi) = 0$. This case is covered by Lemmas 1 and 2, which state that N_d and A_d are empty for all $d \in D$. As a result, the equations from 1 and 2 can be rewritten as:

$$\rho\big((p,\psi),m_1\big) = \bigvee_{d \in D} \Big(\bigwedge_{e \in E_d} \rho\big((p_{d,e},\varepsilon_{d,e}),m_1\big) \Big) \tag{8}$$

$$P[p,\psi](\overline{m}) \Leftrightarrow \bigvee_{d\in D}\Big(\bigwedge_{e\in E_d} P[p_{d,e},\varepsilon_{d,e}](\overline{m})\Big). \tag{9}$$

Since every formula $\varepsilon_{d,e}$ is either an equality or an inequality, it follows from the definition of ρ that $\rho\big((p_{d,e},\varepsilon_{d,e}),m_1\big)$ can only output $(\varnothing,\top)$ or $(\varnothing,\bot)$ and that the couple $(\varnothing,\top)$ is output if and only if $P[p_{d,e},\varepsilon_{d,e}](\overline{m})$ is true.

Since the only successor of a node $(\varnothing,\top)$ is itself, and since $\top$ is in F_ψ, this node generates an accepting branch. A node $(\varnothing,\bot)$ also loops back on itself, but since $\bot$ is not in F_ψ, the resulting branch is not accepting. Thus, the fact that there exists a F_ψ-accepting run $r\big((p,\psi),\overline{m}\big)$ implies that there exists a $d^*\in D$ for which $P[p_{d^*,e},\varepsilon_{d^*,e}](\overline{m})$ holds for all $e\in E_{d^*}$. The right-hand side of (9) holds as a result, and lemma 5 holds in the base case.

Induction step: For some natural $t>0$, if Lemma 5 holds for any formula of depth less than or equal to t, the it holds for depth t.

Let $depth(\psi)=t$. The sets N_d and A_d may not be empty for all $d\in D$. By the definition of a run of ρ, the children of the root (p,ψ) in the accepting run $r\big((p,\psi),\overline{m}\big)$ satisfy the right-hand side of equation (6). Hence, there exists a $d^*\in D$ for which every output $\rho\big((p_{d^*,e},\varepsilon_{d^*,e}),m_1\big)$ and every couple $(p_{d^*,n},\psi_{d^*,n})$ or $(\varnothing,\top_{d,a})$ in (6) is a child of the root. From the base case, we can conclude that $P[p_{d^*,e},\varepsilon_{d^*,e}](\overline{m})$ is true for any $e\in E_{d^*}$. It remains to show that $P[p_{d^*,n},\psi_{d^*,n}](\overline{m}^2)$ is true for any $n\in N_{d^*}$ and every $P[\varnothing,\top_{d,a}](\overline{m}^2)$ is true for any $a\in A_{d^*}$.

We consider each case in turn.

A couple $(p_{d^*,n},\psi_{d^*,n})$ is the root of at least one subrun in $r\big((p,\psi),\overline{m}\big)$ starting at message m_2. Let's denote it by $r\big((p_{d^*,n},\psi_{d^*,n}),\overline{m}^2\big)$. We know from Lemma 1 that $p_{d^*,n}$ assigns every free variable in $\psi_{d^*,n}$. We can also argue that this subrun is $F_{\psi_{d^*,n}}$-accepting since every branch in $r\big((p_{d^*,n},\psi_{d^*,n}),\overline{m}^2\big)$ is an infinite suffix of a branch in $r\big((p,\psi),\overline{m}\big)$. Thus, F_ψ is still visited infinitely often, but the formulas are limited to $sub(\psi_{d^*,n})$ (Lemma 1). As a result, the intersection $F_\psi\cap(sub(\psi_{d^*,n})\cup\{\top\})=F_{\psi_{d^*,n}}$ is visited infinitely often which satisfies the antecedent of Lemma 5.

The rest of the argument is made easy by the induction hypothesis. Indeed, by Lemma 1, a formula $\psi_{d^*,n}$ fits only one of three possible descriptions:

1) $depth(\psi_{d^*,n})<t$

The induction hypothesis can be applied directly to $r\big((p_{d^*,n},\psi_{d^*,n}),\overline{m}^2\big)$. Therefore, $P[p_{d^*,n},\psi_{d^*,n}](\overline{m}^2)$ holds.

2) $depth(\psi_{d^*,n})=t$ and $\psi_{d^*,n}=\mu\ \mathbf{U}\ \eta$ where $\mu,\eta\in sub(\psi)$

According to the case $\mathbf{U}$ of Lemma 4, $r\big((p_{d^*,n},\mu\ \mathbf{U}\ \eta),\overline{m}^2\big)$ implies that:

- $\exists\ j\geq 2$ for which ρ admits an F_η-accepting run with root $(p_{d^*,n},\eta)$;
- $\forall\ 2\leq i<j$, ρ admits an F_μ-accepting run with root $(p_{d^*,n},\mu)$.

Since $depth(\mu)$ and $depth(\eta)$ are less than $depth(\mu\ \mathbf{U}\ \eta) = t$, we can apply the induction hypothesis to the runs listed above:

- $\exists\ j \geq 2$ for which $P[p_{d^*,n}, \eta](\overline{m}^j)$ holds;
- $\forall\ 2 \leq i < j$, $P[p_{d^*,n}, \mu](\overline{m}^i)$ holds.

By definition of the $\mathbf{U}$ operator, $P[p_{d^*,n}, \mu\ \mathbf{U}\ \eta](\overline{m}^2)$ holds.

3) $depth(\psi_{d^*,n}) = t$ and $\psi_{d^*,n} = \mu\,\mathbf{R}\,\eta$ where $\mu, \eta \in sub(\psi)$

The proof is analogous to the one for 2). Upon using Lemma 4 and the induction hypothesis, we get the following possibilities:

- $\forall\ i \geq 2$, $P[p_{d^*,n}, \eta](\overline{m}^i)$holds;

or

- $\exists\ j \geq 2$ for which $P[p_{d^*,n}, \mu \wedge \eta](\overline{m}^j)$ holds;
- $\forall\ 2 \leq i < j$, $P[p_{d^*,n}, \eta](\overline{m}^i)$ holds.

By the definition of the $\mathbf{R}$ operator, $P[p_{d^*,n}, \mu\,\mathbf{R}\,\eta](\overline{m}^2)$ holds in both cases.

We now turn to the set A_{d^*}. Let $P[\varnothing, \top_{d^*,a}](\overline{m})$ be state in A^*_d. Since the run $r\big((p,\psi), \overline{m}\big)$ which generates this state is accepting, $P[\varnothing, \top_{d^*,a}](\overline{m})$ is necessarily of the form $P[\varnothing, \top](\overline{m})$. By the definition of the transition function, this state is reached after the application of the transition function to a formula of the form $\psi = \forall_\pi : \psi'$ for which the $Dom_m(\pi)$ is empty, for some message $\overline{m}'$ and valuation function p. Since $depth(\psi) < t$, we can conclude form the induction hypothesis that $P[p, \psi](\overline{m}')$ holds.

In conclusion, since every term $P[p_{d,e}, \varepsilon_{d,e}](\overline{m})$ and $P[p_{d,n}, \psi_{d,n}](\overline{m}^2)$ in (7) holds for some $d \in D$, it follows that $P[p, \psi](\overline{m})$ holds. □

Proposition 1 *For any LTL-FO$^+$ formula φ devoid of free variables, and any message trace $\overline{m}$, if the automaton A_φ accepts $\overline{m}$, then $\overline{m}$ satisfies φ.*

Proof. An accepting run of A_φ on $\overline{m}$ is, by definition, an F_φ-accepting run of ρ with root $(\varnothing, \varphi)$. By Lemma 5, $r\big((\varnothing, \psi), \overline{m}\big)$ is accepting implies that $P[\varnothing, \varphi](\overline{m})$ is true, which in turns means that "$\overline{m}$ satisfies φ given the free variable assignments in $\varnothing$". □

5.3 $\overline{m} \models \varphi \Rightarrow A_\varphi$ accepts $\overline{m}$

The following lemma states the connection between the semantics of the the temporal operator $\mathbf{R}$ and $\mathbf{U}$ and the tree decomposition of these formulæ. It is analogous to Lemma 4.

Lemma 6 *Let μ and η be LTL-FO$^+$ formulas, and $p : V \to D$ be a partial function that assigns every free variable in μ and η. Let $\overline{m} = m_1, m_2, \ldots$ be a message trace. If:*

- $\forall\ i \geq 1$, *ρ admits an F_η-accepting run $r\big((p,\eta), \overline{m}^i\big)$ with root (p, η),*

then ρ admits an $F_{(\mu \,\mathbf{R}\, \eta)}$-accepting run $r\big((p, \mu \,\mathbf{R}\, \eta), \overline{m}\big)$ with root $(p, \mu \,\mathbf{R}\, \eta)$. Furthermore, if:

- *$\exists\ j \geq 1$ for which ρ admits an $F_{(\mu \wedge \eta)}$-accepting run $r\big((p, \mu \wedge \eta), \overline{m}^j\big)$ with root $(p, \mu \wedge \eta)$ (resp. F_η-accepting run $r\big((p, \eta), \overline{m}^j\big)$ with root (p, η));*
- *$\forall\ 1 \leq i < j$, ρ admits an F_η-accepting run $r\big((p, \eta), \overline{m}^i\big)$ with root (p, η) (resp. F_μ-accepting run $r\big((p, \mu), \overline{m}^i\big)$ with root (p, μ)),*

then ρ admits an $F_{(\mu \,\mathbf{R}\, \eta)}$-accepting run $r\big((p, \mu \,\mathbf{R}\, \eta), \overline{m}\big)$ with root $(p, \mu \,\mathbf{R}\, \eta)$ (resp. $F_{(\mu \ \mathbf{U}\ \eta)}$-accepting run $r\big((p, \mu \ \mathbf{U}\ \eta), \overline{m}\big)$ with root $(p, \mu \ \mathbf{U}\ \eta)$).

Proof. 1) Case " $\forall\ i \geq 1$, ρ admits an F_η-accepting run $r\big((p, \eta), \overline{m}^i\big)$ on $\overline{m}^i$ and with root (p, η)"

Let's suppose that ρ admits an F_η-accepting run $r\big((p, \eta), \overline{m}^i\big)$ for all $i \geq 1$. We must build an $F_{(\mu \,\mathbf{R}\, \eta)}$-accepting run $r\big((p, \mu \,\mathbf{R}\, \eta), \overline{m}\big)$ that includes each one of these runs.

The beginning of any run $r\big((p, \eta), \overline{m}^i\big)$, and of the run $r\big((p, \mu \,\mathbf{R}\, \eta), \overline{m}\big)$, is dictated by the rule

$$\rho\big((p, \mu \,\mathbf{R}\, \eta), m\big) = \rho\big((p, \mu \wedge \eta), m\big) \vee \Big(\rho\big((p, \eta), m\big) \wedge (p, \mu \,\mathbf{R}\, \eta)\Big),$$

which is also applied each time the children of a node $(p, \mu \,\mathbf{R}\, \eta)$ are considered. Since p assigns every free variable in μ and η, we can use Lemma 1 as we did in the proof of Lemma 4 to obtain the rule

$$\rho\big((p, \mu \,\mathbf{R}\, \eta), m\big) = \bigvee_{d \in D_{(\mu \wedge \eta)}} \Big(\bigwedge_{c \in C_d} R_{d,c} \Big) \vee \bigvee_{d \in D_\eta} \Big(\bigwedge_{c \in C_d} R_{d,c} \wedge (p, \mu \,\mathbf{R}\, \eta) \Big). \quad (10)$$

The sets $D_{(\mu \wedge \eta)}$ and D_η encompasses the decomposition of $\rho\big((p, \mu \wedge \eta), m\big)$ and $\rho\big((p, \eta), m\big)$ respectively.

By Definition 4, in a run $r\big((p, \eta), \overline{m}^i\big)$, there exits at least one $d_\eta^i \in D_\eta$ for which any couple $R_{d_\eta^i, c}$ in (10) is a child of the root (if $c \in E_{d_\eta^i}$, the "couple $R_{d_\eta^i, c}$" is the output of $R_{d_\eta^i, c}$). Let's denote the set of all couples $R_{d_\eta^i, c}$ by $R_{d_\eta^i}$. Similarly, for each node $(p, \mu \,\mathbf{R}\, \eta)$ in the run $r\big((p, \mu \,\mathbf{R}\, \eta), \overline{m}\big)$, there must be at least one $d^* \in D_{(\mu \wedge \eta)} \cup D_\eta$ for which any couple in (10) indexed by d^* is a child of $(p, \mu \,\mathbf{R}\, \eta)$.

Let's consider the case where $i = 1$. If we set $d^* = d_\eta^1$ for the root of the node $r\big((p, \mu \,\mathbf{R}\, \eta), \overline{m}\big)$, then (10) implies that the children must at least include the set $R_{d_\eta^1}$ and a node $(p, \mu \,\mathbf{R}\, \eta)$. We choose not to include others nodes, so it remains to define the subruns generated by $R_{d_\eta^1}$ and $(p, \mu \,\mathbf{R}\, \eta)$. The subruns generated by the former are easy because they can be copied from $r\big((p, \eta), \overline{m}^1\big)$. Since this run is F_η-accepting, we know that every branch in a copied subrun visits F_η infinitely often. With the inclusion $F_\eta \subset F_{(\mu \,\mathbf{R}\, \eta)}$, we can also state that every branch visits $F_{(\mu \,\mathbf{R}\, \eta)}$ infinitely often. As for the child node $(p, \mu \,\mathbf{R}\, \eta)$, its subrun is obtained by repeating the previous procedure for $i = 2$ and beyond. We use d_η^i for every $i \geq 2$, and we copy subruns in $r\big((p, \eta), \overline{m}^i\big)$ for $R_{d_\eta^i}$.

The resulting run of ρ on $\overline{m}$ and with root $(p, \mu \,\mathbf{R}\, \eta)$ admits two types of branches. All but one eventually reach a node $R_{d^i_\eta, c}$ for some $i \geq 1$, and thus visit $F_{(\mu \,\mathbf{R}\, \eta)}$ infinitely often. The remaining branch never reaches a node $R_{d^i_\eta, c}$, and thus only visits the node $(p, \mu \,\mathbf{R}\, \eta)$. Fortunately, the formula $\mu \,\mathbf{R}\, \eta$ identifies an accepting state, so this branch also visits $F_{(\mu \,\mathbf{R}\, \eta)}$ infinitely often. These last observations make the resulting run $F_{(\mu \,\mathbf{R}\, \eta)}$-accepting.

2) Case : " $\exists\ j \geq 1$ for which ρ admits an $F_{(\mu \wedge \eta)}$-accepting run $r\big((p, \mu \wedge \eta), \overline{m}^j\big)$ on $\overline{m}^j$ and with root $(p, \mu \wedge \eta)$ (resp. F_η-accepting run $r\big((p, \eta), \overline{m}^j\big)$ on $\overline{m}^j$ and with root (p, η));
$\forall\ 1 \leq i < j$, ρ admits an F_η-accepting run $r\big((p, \eta), \overline{m}^i\big)$ on $\overline{m}^i$ and with root (p, η) (resp. F_μ-accepting run $r\big((p, \mu), \overline{m}^i\big)$ on $\overline{m}^i$ and with root (p, μ))".

Let's now suppose that ρ admits an $F_{(\mu \wedge \eta)}$-accepting run $r\big((p, \mu \wedge \eta), \overline{m}^j\big)$ for some $j \geq 1$. Let's also suppose that for any $1 \leq i < j$, ρ admits an F_η-accepting run $r\big((p, \eta), \overline{m}^i\big)$. As with the first "case", we must build an $F_{(\mu \,\mathbf{R}\, \eta)}$-accepting run $r\big((p, \mu \,\mathbf{R}\, \eta), \overline{m}\big)$ that includes all these runs.

Until message m_j in $\overline{m}$ is reached, we can use the approach described in the first "case" to build a partial run from the root $(p, \mu \,\mathbf{R}\, \eta)$. Thus, for all $1 \leq i < j$, a node $(p, \mu \,\mathbf{R}\, \eta)$ in our partial run, upon reading message m_i, is followed by subruns in $r\big((p, \eta), \overline{m}^i\big)$ and by a node $(p, \mu \,\mathbf{R}\, \eta)$. It remains to define a subrun on the trace $\overline{m}^j$ generated by a node $(p, \mu \,\mathbf{R}\, \eta)$.

By Definition 4, in the run $r\big((p, \mu \wedge \eta), \overline{m}^j\big)$, there is an index $d^j \in D_{(\mu \wedge \eta)}$ for which a couple $R_{d,c}$ in (10) is a child of the root $(p, \mu \wedge \eta)$ if $d = d^j$. Let's denote the set of all couples $R_{d^j, c}$ by R_{d^j}. Not only does R_{d^j} satisfies the right-hand side of (10), the subrun of any couple $R_{d^j, c}$ in $r\big((p, \mu \wedge \eta), \overline{m}^j\big)$ visits $F_{(\mu \wedge \eta)}$ infinitely often. Since $F_{(\mu \wedge \eta)} \subset F_{(\mu \,\mathbf{R}\, \eta)}$, it also visits $F_{(\mu \,\mathbf{R}\, \eta)}$ infinitely often. We hence use the couples in R_{d^j} and their subruns to complete our run on $\overline{m}$.

Any branch β in the resulting run eventually reaches a branch in a copied subrun on $\overline{m}^k$ for some $1 \leq k \leq j$ specific to β. As previously argued, these subbranches visit $F_{(\mu \,\mathbf{R}\, \eta)}$ infinitely often. This makes our run on $\overline{m}$ $F_{(\mu \,\mathbf{R}\, \eta)}$-accepting.

The proof for the subcase $\mu\ \mathbf{U}\ \eta$ is analogous to the one for $\mu \,\mathbf{R}\, \eta$. It is simply a matter of swapping formulas in some of the symbols used. As an example, equation (10) maintains its overall form, but becomes:

$$\rho\big((p, \mu\ \mathbf{U}\ \eta), m\big) = \bigvee_{d \in D_\eta} \Big(\bigwedge_{c \in C_d} R_{d,c} \Big) \vee \bigvee_{d \in D_\mu} \Big(\bigwedge_{c \in C_d} R_{d,c} \wedge (p, \mu\ \mathbf{U}\ \eta) \Big).$$

The sets F_μ and F_η are obviously included in the set $F_{(\mu\ \mathbf{U}\ \eta)}$. □

Lemma 7 *For any LTL-FO$^+$ formula ψ and any partial function $p : V \to I$ that assigns every free variable in ψ, if the predicate $P[p, \psi](\overline{m})$ is true for a trace $\overline{m}$, then ρ admits an F_ψ-accepting run $r\big((p, \psi), \overline{m}\big)$ with root (p, ψ).*

Proof. As with Lemma 5, we proceed by strong induction on the temporal depth of LTL-FO$^+$ formulas. Our approach is based on the decompositions of $\rho\big((p,\psi),m_1\big)$ and $P[p,\psi](\overline{m})$, which are given by Lemmas 1 and 2 respectively:

$$\rho\big((p,\psi),m_1\big) = \bigvee_{d\in D}\Big(\bigwedge_{e\in E_d}\rho\big((p_{d,e},\varepsilon_{d,e}),m_1\big)\wedge\bigwedge_{n\in N_d}(p_{d,n},\psi_{d,n})\wedge\bigwedge_{a\in A_d}(\varnothing,\top_{d,a})\Big); \tag{6}$$

$P[p,\psi](\overline{m}) \Leftrightarrow$

$$\bigvee_{d\in D}\Big(\bigwedge_{e\in E_d}P[p_{d,e},\varepsilon_{d,e}](\overline{m})\wedge\bigwedge_{n\in N_d}P[p_{d,n},\psi_{d,n}](\overline{m}^2)\wedge\bigwedge_{a\in A_d}P[\varnothing,\top_{d,a}](\overline{m}^2)\Big). \tag{7}$$

For any $d\in D$, $e\in E_d$, $n\in N_d$ and $a\in A_d$, the terms $p_{d,e}$, $p_{d,n}$, $\varepsilon_{d,e}$, $\psi_{d,n}$ and $\top_{d,a}$ in (6) can be assumed to be identical to their counterpart in (7) due to Lemma 3.

Base case: Lemma 7 holds for any formula of depth 0.

Suppose $depth(\psi)=0$. We know from Lemmas 1 and 2 that the sets N_d and A_d in this special case, are empty for every $d\in D$. Thus, we can rewrite equations (6) and (7) as follows:

$$\rho\big((p,\psi),m_1\big) = \bigvee_{d\in D}\Big(\bigwedge_{e\in E_d}\rho\big((p_{d,e},\varepsilon_{d,e}),m_1\big)\Big); \tag{8}$$

$$P[p,\psi](\overline{m}) \leftrightarrow \bigvee_{d\in D}\Big(\bigwedge_{e\in E_d}P[p_{d,e},\varepsilon_{d,e}](\overline{m})\Big). \tag{9}$$

If the predicate $P[p,\psi](\overline{m})$ holds, then there exists a $d^*\in D$ for which $P[p_{d^*,e},\varepsilon_{d^*,e}](\overline{m})$ is true for every $e\in E_{d^*}$. It follows from the definition of ρ that the output of a term $\rho\big((p_{d,e},\varepsilon_{d,e}),m_1\big)$ in (8) is $(\varnothing,\top)$ for $d=d^*$. Hence, this couple satisfies the right-hand side of (8) and can thus be the only child of the root (p,ψ) in our run. Since the only successor of a node $(\varnothing,\top)$ is itself (for any input message), the resulting run has a single branch that visits the node $(\varnothing,\top)$ infinitely often. Since the state $\top$ is in F_ψ, this run on $\overline{m}$ is F_ψ-accepting.

Induction step: For some natural $t>0$, if Lemma 7 holds for any formula of depth less than t, then it also holds for any formula of depth t.

Suppose $depth(\psi)=t$. The sets N_d and A_d may not be empty for all $d\in D$, so equations (6) and (7) must be used. If $P[p,\psi](\overline{m})$ holds, then there exists a $d^*\in D$ for which the predicates $P[p_{d^*,e},\varepsilon_{d^*,e}](\overline{m})$, $P[p_{d^*,n},\psi_{d^*,n}](\overline{m}^2)$ and $P[\varnothing,\top_{d,a}](\overline{m}^2)$ in (7) hold for any $e\in E_{d^*}$, $n\in N_{d^*}$ and $a\in A_{d^*}$.

The accepting run for (p,ψ) will include every couple indexed by d^*, in E_{d^*}, N_{d^*} and A_{d^*} . The base case already shows that any term $\rho\big((p_{d,e},\varepsilon_{d,e}),m_1\big)$ indexed by d^* in (6) outputs the accepting state $(\varnothing,\top)$. If N_{d^*} and A_{d^*} are not empty, we must include their content in the children of the root (p,ψ) in the accepting run.

We also include any couple $(p_{d,n}, \psi_{d,n})$ indexed by d^* in (6) in order to satisfy the right-hand side of this equation. We know that the branch generated by the child node $(\varnothing, \top)$ visits F_ψ infinitely often. It remains to define an F_ψ-accepting subrun on the trace $\overline{m}^2$ for every child node $(p_{d^*,n}, \psi_{d^*,n})$. It remains to show that the elements of these sets also generate an accepting run.

We consider first the elements of N_{d^*}.

Since every formula $\psi_{d^*,n}$ is a subformula of ψ (Lemma 1), the inclusion $F_{\psi_{d^*,n}} \subseteq F_\psi$ holds. Therefore, for any child node $(p_{d^*,n}, \psi_{d^*,n})$, it suffices to show that its subrun is $F_{\psi_{d^*,n}}$-accepting. Lemma 1 also tells us that $p_{d^*,n}$ assigns every free variable in $\psi_{d^*,n}$ for any $n \in N_{d^*}$. As such, every predicate $P[p_{d^*,n}, \psi_{d^*,n}](\overline{m}^2)$ in (7) satisfies the condition of Lemma 7.

The remainder of the argument is made easy by the induction hypothesis. Indeed, by Lemma 1, a formula $\psi_{d^*,n}$ fits only one of three possible cases:

1) $depth(\psi_{d^*,n}) < t$

The induction hypothesis directly applies to $P[p_{d^*,n}, \psi_{d^*,n}](\overline{m}^2)$. Thus, ρ admits an $F_{\psi_{d^*,n}}$-accepting run on $\overline{m}^2$ and with root $(p_{d^*,n}, \psi_{d^*,n})$ as desired. As was the case for lemma 5, this case also implies the validity of any formula in A_{d^*}.

2) $depth(\psi_{d^*,n}) = t$ and $\psi_{d^*,n} = \mu \ \mathbf{U} \ \eta$ where $\mu, \eta \in sub(\psi)$

We simply follow the proof of Lemma 5 for this case, but in reverse. First, if $P[p_{d^*,n}, \mu \ \mathbf{U} \ \eta](\overline{m}^2)$ holds, then by the definition of $\mathbf{U}$, for some $j \geq 2$, $P[p_{d^*,n}, \eta](\overline{m}^j)$ holds and for every $2 \leq i < j$, $P[p_{d^*,n}, \mu](\overline{m}^i)$ also holds. Next, because the temporal depths of μ and η are less than t, the induction hypothesis applies. Hence, ρ admits an F_η-accepting run on $\overline{m}^j$ and F_μ-accepting runs on $\overline{m}^i$ for every $2 \leq i < j$. Finally, by Lemma 6, ρ admits an $F_{(\mu \ \mathbf{U} \ \eta)}$-accepting run on $\overline{m}^2$ as desired.

3) $depth(\psi_{d^*,n}) = t$ and $\psi_{d^*,n} = \mu \, \mathbf{R} \, \eta$ where $\mu, \eta \in sub(\psi)$

As with(2), we follow the proof of Lemma 5 for the current case in reverse. Note though that if $P[p_{d^*,n}, \mu \, \mathbf{R} \, \eta](\overline{m}^2)$ holds, two possibilities arise:

- $\forall \ i \geq 2$, $P[p_{d^*,n}, \eta](\overline{m}^i)$ is true;
- $\exists \ j \geq 2$ for which $P[p_{d^*,n}, \mu \wedge \eta](\overline{m}^j)$ is true, and $P[p_{d^*,n}, \eta](\overline{m}^i)$ is true $\forall \ 2 \leq i < j$.

Since the depth of $\mu \wedge \eta$, is less than t, the induction hypothesis applies for both possibilities. □

Proposition 2 *For any LTL-FO$^+$ formula φ devoid of free variables, and let $\overline{m}$ be a message trace, if $\overline{m}$ satisfies φ, then the automaton A_φ accepts $\overline{m}$.*

Proof. All variables in φ are bound by quantifiers ($\exists$ or $\forall$), so the statement "$\overline{m}$ satisfies φ" is represented by $P[\varnothing, \varphi](\overline{m})$. It follows from Lemma 7 that there exists an F_φ-accepting run of ρ on $\overline{m}$ with root $(\varnothing, \varphi)$, which, by definition, is an accepting run of A_φ on $\overline{m}$. □

Theorem 2 (adapted from Vardi) *Given any LTL-FO$^+$ formula φ, one can build a modified alternating Büchi automaton $A_\varphi = (\Sigma, V, I, S, s^0, \rho, F)$, where $\Sigma = M$ and $|S|$ is in $O(|\varphi|)$, such that the language recognized by A_φ is exactly the set of message traces satisfying the formula φ.*

Proof. Immediate from the definition of A_φ and Propositions 1 and 2. □

6 Conclusion and Future Works

In this paper, we propose a new type of finite alternating automata which recognizes LTL-FO$^+$ formulæ an show the process of constructing such an automaton from an LTL-FO$^+$ formula. Our automaton allows for formulæ in the highly expressive logic, LTL-FO$^+$ formal logics to be easily stated in a concise and easy to understand formalism. We are currently developing and implementing a verification algorithm that will allow this new tool to be put to practical use.

References

1. Barringer, H., Falcone, Y., Havelund, K., Reger, G., Rydeheard, D.E.: Quantified event automata: Towards expressive and efficient runtime monitors. In: Giannakopoulou, D., Méry, D. (eds.) FM 2012: Formal Methods - 18th International Symposium, Paris, France, August 27-31, 2012. Proceedings. Lecture Notes in Computer Science, vol. 7436, pp. 68–84. Springer (2012)
2. Büchi, J.: On a decision method in restricted second order arithmetic. In: Logic, methodology and philosophy of science, pp. 1–11 (1962)
3. Clarke, E.M., Grumberg, O., Peled, D.A.: Model Checking. MIT Press (2000)
4. Colombo, C., Pace, G.J., Schneider, G.: Larva — safer monitoring of real-time java programs. In: Seventh IEEE International Conference on Software Engineering and Formal Methods (SEFM). pp. 33–37. IEEE Computer Society (Nov 2009)
5. Fu, X., Bultan, T., Su, J.: Analysis of interacting BPEL web services. In: Feldman, S.I., Uretsky, M., Najork, M., Wills, C.E. (eds.) WWW. pp. 621–630. ACM (2004)
6. Hallé, S., Villemaire, R.: Runtime enforcement of web service message contracts with data. IEEE Trans. Services Computing 5(2), 192–206 (2012), http://dx.doi.org/10.1109/TSC.2011.10
7. Khoury, R., Hallé, S., Waldmann, O.: Execution trace analysis using LTL-FO+. In: 7th International Symposium On Leveraging Applications of Formal Methods, Verification and Validation (IsoLa 16), Corfu, Greece (2016)
8. Nakajima, S.: Lightweight formal analysis of web service flows. Progress in Informatics (2), 57–76 (2005)
9. Robinson, W.: A requirements monitoring framework for enterprise systems. Requir. Eng. 11(1), 17–41 (2006)
10. Vardi, M., Wolper, P.: Reasoning about infinite computations. Inf. Comput. 115(1), 1–37 (Nov 1994), http://dx.doi.org/10.1006/inco.1994.1092